

# **2022 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS 2022)**

**Singapore  
22-24 May 2022**



**IEEE Catalog Number: CFP22PER-POD  
ISBN: 978-1-6654-5955-6**

**Copyright © 2022 by the Institute of Electrical and Electronics Engineers, Inc.  
All Rights Reserved**

*Copyright and Reprint Permissions:* Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

***\*\*\* This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP22PER-POD
ISBN (Print-On-Demand):	978-1-6654-5955-6
ISBN (Online):	978-1-6654-5954-9

**Additional Copies of This Publication Are Available From:**

Curran Associates, Inc  
57 Morehouse Lane  
Red Hook, NY 12571 USA  
Phone: (845) 758-0400  
Fax: (845) 758-2633  
E-mail: [curran@proceedings.com](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

CURRAN ASSOCIATES INC.  
**proceedings**  
.com

# 2022 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS) **ISPASS 2022**

## Table of Contents

Message from the General Chair .....	xi
Message from the Program Chair .....	xiii
Organizing Committee .....	xiv
Program Committee .....	xvi
Steering Committee .....	xviii
Keynotes .....	xix
Sponsors .....	xxi

### Paper Session 1: Best Paper Nominations

Profiling Intel Graphics Architecture with Long Instruction Traces .....	1
<i>Konstantin Levit-Gurevich (Intel Corporation, Israel), Alex Skaletsky (Intel Corporation, Israel), Michael Berezalsky (Intel Corporation, Israel), Yulia Kuznetcova (Intel Corporation, Israel), and Hila Yakov (Intel Corporation, Israel)</i>	
Performance Analysis and Optimization with Little’s Law .....	12
<i>Sanyam Mehta (HPE, USA)</i>	
The Indigo Program-Verification Microbenchmark Suite of Irregular Parallel Code Patterns .....	24
<i>Yiqian Liu (Texas State University, USA), Noushin Azami (Texas State University, USA), Corbin Walters (Texas State University, USA), and Martin Burtscher (Texas State University, USA)</i>	
gpuFI-4: A Microarchitecture-Level Framework for Assessing the Cross-Layer Resilience of Nvidia GPUs .....	35
<i>Dimitris Sartzetakis (University of Athens, Greece), George Papadimitriou (University of Athens, Greece), and Dimitris Gizopoulos (University of Athens, Greece)</i>	

### Paper Session 2: Simulation Technologies and Performance Analysis

Distilling the Real Cost of Production Garbage Collectors .....	46
<i>Zixian Cai (Australian National University, Australia), Stephen M. Blackburn (Australian National University, Australia, and Google, USA), Michael D. Bond (Ohio State University, USA), and Martin Maas (Google, USA)</i>	

Scale-Model Architectural Simulation .....	58
<i>Wenjie Liu (Ghent University, Belgium), Wim Heirman (Intel Corporation, Belgium), Stijn Eyerman (Intel Corporation, Belgium), Shoaib Akram (Australian National University, Australia), and Lieven Eeckhout (Ghent University, Belgium)</i>	
MEGsim: A Novel Methodology for Efficient Simulation of Graphics Workloads in GPUs .....	69
<i>Jorge Ortiz (University of Murcia, Spain), D. Corbalán-Navarro (University of Murcia, Spain), Juan L. Aragón (University of Murcia, Spain), and Antonio González (Polytechnic University of Catalonia, Spain)</i>	
MARTA: Multi-Configuration Assembly pRofiler and Toolkit for Performance Analysis .....	79
<i>Marcos Horro (CITIC, Universidade da Coruña, Spain), Louis-Noël Pouchet (Colorado State University, USA), Gabriel Rodríguez (CITIC, Universidade da Coruña, Spain), and Juan Touriño (CITIC, Universidade da Coruña, Spain)</i>	
Left-Shifter: A pre-Silicon Framework for Usage Model Based Performance Verification of the PCIe Interface in Server Processor System on Chips .....	90
<i>Tessil Thomas (Central Engineering - Systems, Arm® Ltd., United Kingdom), Bharath Venkatasubramanian (Central Engineering - Systems, Arm® Ltd., United Kingdom), Dinesh Sthapit (Central Engineering - Systems, Arm® Ltd., United Kingdom), Christopher Gray (Central Engineering - Systems, Arm® Ltd., United Kingdom), Atresh Gummadavelly (Central Engineering - Systems, Arm® Ltd., United Kingdom), Janick Bergeron (Solutions Group, Synopsys® Inc, USA), Pankaj Mehta (Solutions Group, Synopsys® Inc, India), and Prabu Thangamuthu (Solutions Group, Synopsys® Inc, India)</i>	

### **Paper Session 3: Code Generation and Optimization**

Fourst: A Code Generator for FFT-Based Fast Stencil Computations .....	99
<i>Zafar Ahmad (Stony Brook University, USA), Mohammad Mahdi Javanmard (Meta Platforms, Inc., USA), Gregory Croisdale (The University of Tennessee, Knoxville, USA), Aaron Gregory (Stony Brook University, USA), Pramod Ganapathi (Stony Brook University, USA), Louis-Noël Pouchet (Colorado State University, USA), and Rezaul Chowdhury (Stony Brook University, USA)</i>	
Flexible Binary Instrumentation Framework to Profile Code Running on Intel GPUs .....	109
<i>Alex Skaletsky (Intel Corporation, Israel), Konstantin Levit-Gurevich (Intel Corporation, Israel), Michael Berezalsky (Intel Corporation, Israel), Yulia Kuznetcova (Intel Corporation, Israel), and Hila Yakov (Intel Corporation, Israel)</i>	
POSET-RL: Phase Ordering for Optimizing Size and Execution Time Using Reinforcement Learning .....	121
<i>Shalini Jain (Indian Institute of Technology Hyderabad, India), Yashas Andaluri (Indian Institute of Technology Hyderabad, India), S. VenkataKeerthy (Indian Institute of Technology Hyderabad, India), and Ramakrishna Upadrasta (Indian Institute of Technology Hyderabad, India)</i>	

## Poster Session 1

XFeatur: Hardware Feature Extraction for DNN Auto-Tuning .....	132
<i>Jorge Sierra Acosta (Universitat Politècnica de Catalunya, Spain), Andreas Diavastos (Universitat Politècnica de Catalunya, Spain), and Antonio Gonzalez (Universitat Politècnica de Catalunya, Spain)</i>	
SGXGauge: A Comprehensive Benchmark Suite for Intel SGX .....	135
<i>Sandeep Kumar (Indian Institute of Technology, Delhi, India), Abhisek Panda (Indian Institute of Technology, Delhi, India), and Smruti R. Sarangi (Indian Institute of Technology, Delhi, India)</i>	
SAPCo Sort: Optimizing Degree-Ordering for Power-Law Graphs .....	138
<i>Mohsen Koochi Esfahani (Queen's University Belfast, UK), Peter Kilpatrick (Queen's University Belfast, UK), and Hans Vandierendonck (Queen's University Belfast, UK)</i>	
TILE-SIM: A Systematic Approach to Systolic Array-Based Accelerator Evaluation .....	141
<i>Yuhang Li (National University of Defense Technology, China), Mei Wen (National University of Defense Technology, China), Jiawei Fei (National University of Defense Technology, China), Junzhong Shen (National University of Defense Technology, China), and Yasong Cao (National University of Defense Technology, China)</i>	
FASE: A Fast, Accurate and Seamless Emulator for Custom Numerical Formats .....	144
<i>John Osorio (Barcelona Supercomputing Center, Universitat Politècnica de Catalunya, Spain), Adrià Armejach (Barcelona Supercomputing Center, Universitat Politècnica de Catalunya, Spain), Eric Petit (Intel, USA), Greg Henry (Intel, USA), and Marc Casas (Barcelona Supercomputing Center, Universitat Politècnica de Catalunya, Spain)</i>	
Meterstick: Benchmarking Performance Variability in Cloud and Self-Hosted Minecraft-Like Games .....	147
<i>Jerrit Eickhoff (Delft University of Technology, Netherlands), Jesse Donkeroliet (Vrije Universiteit Amsterdam, Netherlands), and Alexandru Iosup (Vrije Universiteit Amsterdam, Netherlands)</i>	
Simulating Noisy Channels in DNA Storage .....	150
<i>Mayank Keoliya (National University of Singapore, Singapore), Puru Sharma (National University of Singapore, Singapore), and Djordje Jevdjic (National University of Singapore, Singapore)</i>	
OS-Level Implications of Using DRAM Caches in Memory Disaggregation .....	153
<i>Bin Gao (National University of Singapore, Singapore), Hao Wei Tee (National University of Singapore, Singapore), Alireza Sanaee (Queen Mary University of London, United Kingdom), Boon Jun Soh (National University of Singapore, Singapore), and Djordje Jevdjic (National University of Singapore, Singapore)</i>	
VIPP: Validation-Included Precision-Parametric N-Body Benchmark Suite .....	156
<i>Shigeyuki Sato (The University of Tokyo, Japan), Kota Iizuka (Fixstars Corp., Japan), Naoki Yoshifuji (Fixstars Corp., Japan), and Masaki Natsume (Fixstars Corp., Japan)</i>	

High-Performance Deployment of Text Detection Model: Compression and Hardware Platform Considerations .....	159
<i>Nupur Sumeet (Tata Consultancy Services Research, India), Karan Rawat (Tata Consultancy Services Research, India), and Manoj Nambiar (Tata Consultancy Services Research, India)</i>	

## Paper Session 4: Benchmarks and Characterization of Emerging Workloads, Scheduling and Reliability

Roofline Model for UAVs: A Bottleneck Analysis Tool for Onboard Compute Characterization of Autonomous Unmanned Aerial Vehicles .....	162
<i>Srivatsan Krishnan (Harvard University, USA), Zishen Wan (Harvard University, USA), Kshitij Bhardwaj (Lawrence Livermore National Lab, USA), Ninad Jadhav (Harvard University, USA), Aleksandra Faust (Google Brain Research, USA), and Vijay Janapa Reddi (Harvard University, USA)</i>	
RTRBench: A Benchmark Suite for Real-Time Robotics .....	175
<i>Mohammad Bakhshalipour (Carnegie Mellon University, USA), Maxim Likhachev (Carnegie Mellon University, USA), and Phillip B. Gibbons (Carnegie Mellon University, USA)</i>	
Characterization of MPC-Based Private Inference for Transformer-Based Models .....	187
<i>Yongqin Wang (University of Southern California, USA), G. Edward Suh (Meta AI, USA), Wenjie Xiong (Meta AI, USA), Benjamin Lefaudeux (Meta AI, USA), Brian Knott (Meta AI, USA), Murali Annavaram (University of Southern California, USA), and Hsien-Hsin S. Lee (Meta AI, USA)</i>	
Spatiotemporal Strategies for Long-Term FPGA Resource Management .....	198
<i>Atefeh Mehrabi (Duke University, USA), Daniel J. Sorin (Duke University, USA), and Benjamin C. Lee (University of Pennsylvania, USA)</i>	
Eris: Fault Injection and Tracking Framework for Reliability Analysis of Open-Source Hardware .....	210
<i>Shubham Nema (North Carolina State University, USA), Justin Kirschner (North Carolina State University, USA), Debpratim Adak (North Carolina State University, USA), Sapan Agarwal (Sandia National Laboratories, USA), Ben Feinberg (Sandia National Laboratories, USA), Arun F. Rodrigues (Sandia National Laboratories, USA), Matthew J. Marinella (Arizona State University, USA), and Amro Awad (North Carolina State University, USA)</i>	

## Poster Session 2

Understanding Data Compression in Warehouse-Scale Datacenter Services .....	221
<i>Geonhwa Jeong (Georgia Institute of Technology, USA), Bikash Sharma (Meta Inc., USA), Nick Terrell (Meta Inc., USA), Abhishek Dhanotia (Meta Inc., USA), Zhiwei Zhao (Meta Inc., USA), Niket Agarwal (Meta Inc., USA), Arun Kejariwal (Meta Inc., USA), and Tushar Krishna (Georgia Institute of Technology, USA)</i>	

LoopIn: A Loop-Based Simulation Sampling Mechanism .....	224
<i>Uday Kumar Reddy Vengalam (University of Rochester, USA), Anshujit Sharma (University of Rochester, USA), and Michael Huang (University of Rochester, USA)</i>	
Building a Performance Model for Deep Learning Recommendation Model Training on GPUs .....	227
<i>Zhongyi Lin (University of California, Davis, USA), Louis Feng (Meta Platforms, Inc, USA), Ehsan K. Ardestani (Meta Platforms, Inc, USA), Jaewon Lee (Meta Platforms, Inc, USA), John Lundell (Meta Platforms, Inc, USA), Changkyu Kim (Meta Platforms, Inc, USA), Arun Kejariwal (Meta Platforms, Inc, USA), and John D. Owens (University of California, Davis, USA)</i>	
Advancing Near-Data Processing with Precise Exceptions and Efficient Data Fetching .....	230
<i>Sairo Santos (Federal University of Paraná – Curitiba, Brazil; Federal Rural University of the Semi-arid – Angicos, Brazil), Tiago R. Kepe (Federal University of Paraná – Curitiba, Brazil), Francis B. Moreira (Federal University of Rio Grande do Sul - Porto Alegre, Brazil), Paulo C. Santos (Federal University of Rio Grande do Sul - Porto Alegre, Brazil), and Marco A. Z. Alves (Federal University of Paraná – Curitiba, Brazil)</i>	
Profiling an Architectural Simulator .....	233
<i>Nedasadat Taheri (University of Kansas, USA), Alexander Manley (University of Kansas, USA), Ahmni R Pang (University of Kansas, USA), and Mohammad Alian (University of Kansas, USA)</i>	
Benchmarking Test-Time Unsupervised Deep Neural Network Adaptation on Edge Devices .....	236
<i>Kshitij Bhardwaj (Lawrence Livermore National Laboratory, USA), James Diffenderfer (Lawrence Livermore National Laboratory, USA), Bhavya Kailkhura (Lawrence Livermore National Laboratory, USA), and Maya Gokhale (Lawrence Livermore National Laboratory, USA)</i>	
GPUCalorie: Floorplan Estimation for GPU Thermal Evaluation .....	239
<i>Marcus Chow (University of California Riverside, USA), Ali Jahanshahi (University of California Riverside, USA), Ana Cardenas Beltran (University of California Riverside, USA), Sheldon Tan (University of California Riverside, USA), and Daniel Wong (University of California Riverside, USA)</i>	
ARBench: Augmented Reality Benchmark For Mobile Devices .....	242
<i>Sofiane Chetoui (Brown University, USA), Rahul Shahi (Brown University, USA), Seif Abdelaziz (Cupertino High School, USA), Abhinav Golas (Meta Platforms, USA), Farrukh Hijaz (Meta Platforms, USA), and Sherief Reda (Brown University, USA)</i>	
Cross-Level Characterization of Program Behavior .....	245
<i>Li Tang (Los Alamos National Laboratory, USA) and Scott Pakin (Los Alamos National Laboratory, USA)</i>	
A SIMT Analyzer for Multi-Threaded CPU Applications .....	248
<i>Ahmad Alawneh (Purdue University, USA), Mahmoud Khairy (Purdue University, USA), and Timothy G. Rogers (Purdue University, USA)</i>	

Microarchitectural Performance Evaluation of AV1 Video Encoding Workloads .....	251
<i>Steffen Jensen (The University of Texas at Austin), Jaekyu Lee (Arm Research), Dam Sunwoo (Arm Research), Matthew J. Horsnell (Arm Research), and Lizy K. John (The University of Texas at Austin)</i>	

## Paper Session 5: Machine Learning Systems, Acceleration and Applications

Ruby: Improving Hardware Efficiency for Tensor Algebra Accelerators Through Imperfect Factorization .....	254
<i>Mark Horeni (University of Notre Dame, USA), Pooria Taheri (University of Notre Dame, USA), Po-An Tsai (NVIDIA, USA), Angshuman Parashar (NVIDIA, USA), Joel Emer (MIT/NVIDIA, USA), and Siddharth Joshi (University of Notre Dame, USA)</i>	
Pareto Rank Surrogate Model for Hardware-Aware Neural Architecture Search .....	267
<i>Hadjer Benmeziane (Université Polytechnique Hauts-de-France, France), Smail Niar (Université Polytechnique Hauts-de-France, France), Hamza Ouarnoughi (Université Polytechnique Hauts-de-France, France), and Kaoutar El Maghraoui (IBM T. J. Watson Research Center, USA)</i>	
Learning A Continuous and Reconstructible Latent Space for Hardware Accelerator Design .....	277
<i>Qijing Huang (NVIDIA, USA), Charles Hong (University of California, Berkeley, USA), John Wawrzyniek (University of California, Berkeley, USA), Mahesh Subedar (Intel Labs, USA), and Yakun Sophia Shao (University of California, Berkeley, USA)</i>	
Bifrost: End-to-End Evaluation and Optimization of Reconfigurable DNN Accelerators .....	288
<i>Axel Stjerngren (University of Glasgow, United Kingdom), Perry Gibson (University of Glasgow, United Kingdom), and José Cano (University of Glasgow, United Kingdom)</i>	

## Paper Session 6: Memory Technologies and Architectures

PCMCsim: An Accurate Phase-Change Memory Controller Simulator and its Performance Analysis....	300
<i>Hyokeun Lee (Seoul National University, South Korea), Hyungsuk Kim (Seoul National University, South Korea), Seokbo Shim (SK Hynix, South Korea), Seungyong Lee (Seoul National University, South Korea), Dosun Hong (SK Hynix, South Korea), Hyuk-Jae Lee (Seoul National University, South Korea), and Hyun Kim (Seoul National University of Science and Technology, South Korea)</i>	
Address Translation Conscious Caching and Prefetching for High Performance Cache Hierarchy...	311
<i>Vasudha Vasudha (Indian Institute of Technology Kanpur, India) and Biswabandan Panda (Indian Institute of Technology Kanpur, India)</i>	
DRAM Bandwidth and Latency Stacks: Visualizing DRAM Bottlenecks .....	322
<i>Stijn Eyerman (Intel Corporation, Belgium), Wim Heirman (Intel Corporation, Belgium), and Ibrahim Hur (Intel Corporation, USA)</i>	

Author Index .....	333
--------------------	-----